

The Bioinformatics Lab

Challenge

Network filesystems and Grid Computing

Jonas Reeb

25.06.2013

Layout

NFS

Samba

SSHFS

Sun Grid Engine

Installation

```
apt-get install nfs-kernel-server  
#nfs-common in dependencies
```

```
[warn] Not starting NFS kernel daemon: no exports. *** (warning).
```

/etc/exports |

```
/home *(ro,root_squash,all_squash)
/home jonas.tbl(rw,no_root_squash)
/home 192.168.16.34(rw,no_root_squash)
/home 127.0.0.1(rw,no_root_squash)
```

- ▶ `fsid=0` not required anymore, default root is /
⇒ NFS v2/3/4 all see the same path

/etc/exports II

- ▶ **No space** between host and options

```
/home jonas.tbl (rw)
```

- ▶ gives `ro` to `jonas.tbl` but `rw` to the world!

/etc/exports options

- ▶ ro
Mount read-only, no changes possible even for root
- ▶ (root|all)_squash
Client root or every user is demoted to privileges of user `nfsnobody`
- ▶ (no_)subtree_check
(Do not) check on every access whether file resides in exported area

Start and mount

```
/etc/init.d/nfs-kernel-server restart  
#(re)-start NFS server  
exportfs -vra  
#Upon every change of /etc/exports
```

```
mount -t nfs4 192.168.16.4:/home  
/mnt/jonas-home #NFSv4  
mount -t nfs 192.168.16.4:/home  
/mnt/jonas-home #NFSv2/3
```

/etc/fstab

```
jonas.tbl:/home /mnt/jonas-home nfs4  
noauto 0 0
```

Layout

NFS

Samba

SSHFS

Sun Grid Engine

Setup

Install

```
apt-get install samba samba-doc smbclient
```

Log files are in `/var/log/samba/log.*`

`/etc/samba/smb.conf`

```
unix password sync = no #Seperate passwds  
workgroup = TBL #unless set during install  
log level = 10 all:10 #Spammy but helpful
```

Change/set password for user `jonas`

```
smbpasswd -a jonas
```

Create share

/etc/samba/smb.conf

```
security = users
[...]
[clipboard]
    path = /srv/samba/clipboard
    writeable = yes                #Default: no
    browseable = yes              #Show in list of avail. shares (default)
    valid users = jonas
    force create mode = 0777      #Default: 0744
    force directory mode = 0777  #Default: 0000
```

Test for correct syntax

```
testparm /etc/samba/smb.conf
```

Restart daemons

```
/etc/init.d/samba restart
```

Access shares

- ▶ CLI:

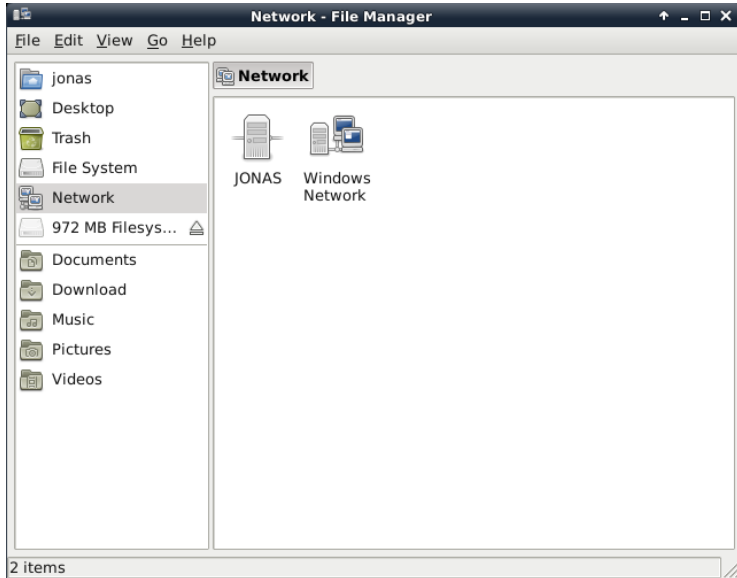
```
smbclient -U jonas -L jonas.tbl  
#List all shares
```

- ▶ X:

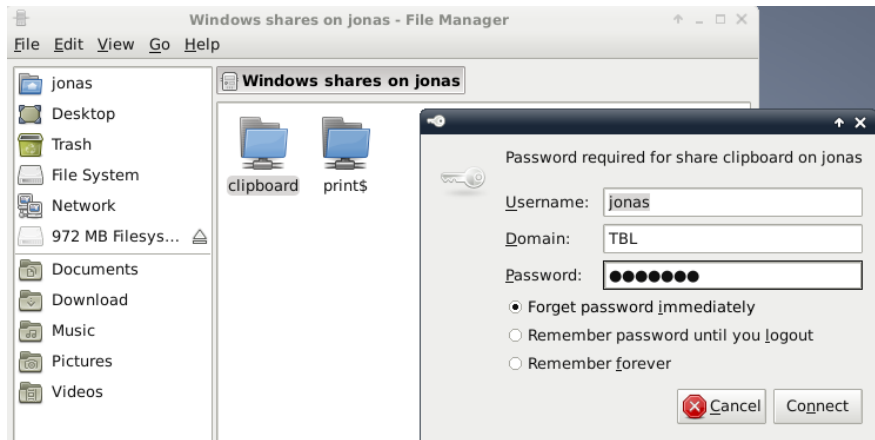
Filebrowser Thunar (Xfce default) requires additional packages

```
apt-get install gvfs gvfs-backends
```

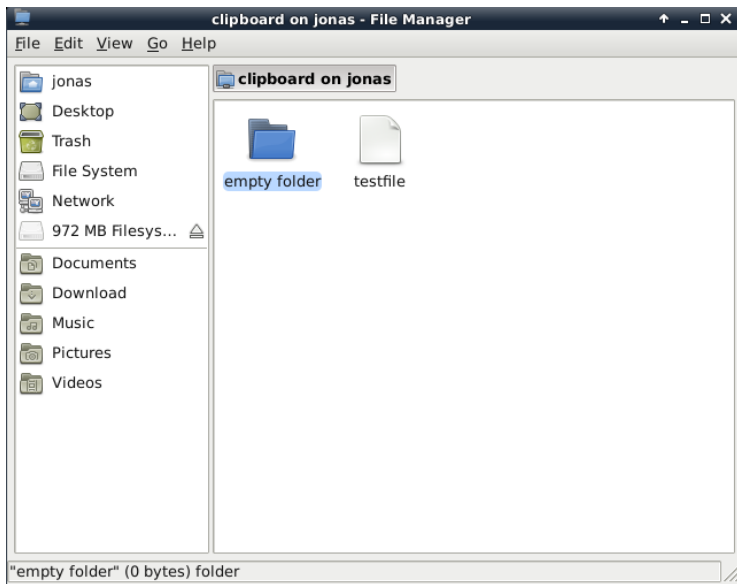
View in Thunar I



View in Thunar II



View in Thunar III



Notes

- ▶ Local permissions weigh stronger than Samba
`chmod -R 777 /srv/samba/clipboard`
- ▶ SWAT can be used for web administration

Advanced: Mounting on L2 VM

Use `mount` with `cifs` protocol

```
apt-get install cifs-utils  
mkdir clipl1  
mount -t cifs //192.168.16.4/clipboard  
./clipl1 -o user=jonas
```

Advanced: Idapsam - Schema I

Get the schema file

```
apt-get install samba-doc  
cd /etc/share/doc/samba-doc/examples/LDAP/  
gunzip samba.schema.gz  
cp samba.schema /etc/ldap/schema
```

Create /etc/ldap/schema/schema_convert.conf

```
include /etc/ldap/schema/core.schema  
include /etc/ldap/schema/cosine.schema  
include /etc/ldap/schema/inetorgperson.  
    schema  
include /etc/ldap/schema/samba.schema
```

Advanced: Idapsam - Schema II

Convert schemas to .ldif using slapcat

```
cd /etc/ldap/schema/  
mkdir /tmp/ldif_out  
slapcat -f schema_convert.conf  
    -F /tmp/ldif_out -n0  
cd /tmp/ldif_out/cn\=config/cn\=schema/  
cp cn\=\{* /etc/ldap/schema/
```

Advanced: Idapsam - Schema III

- ▶ Rename to `samba.ldif`
- ▶ Adjust `dn:` and `cn:` line
- ▶ Delete last seven lines

Add schema

```
ldapadd -Y EXTERNAL -H ldapi:///
-f samba.ldif
```

<http://www.samba.org/samba/docs/man/Samba-HOWTO-Collection/passdb.html>,
http://wiki.ubuntuusers.de/OpenLDAP_ab_Precise,
<https://help.ubuntu.com/12.04/serverguide/samba-ldap.html>

Advanced: ldapsam - smb.conf I

Adjust smb.conf (see man smb.conf) for ldap auth

```
#passdb backend = tdbsam
passdb backend = ldapsam

ldapsam:trusted=yes
ldapsam:editposix=yes

ldap admin dn = cn=admin,dc=tbl
ldap delete dn = yes
ldap group suffix = ou=groups
ldap idmap suffix = ou=idmap
ldap machine suffix = ou=machines
ldap user suffix = ou=people
ldap suffix = dc=tbl

idmap backend = ldap:"ldap://192.168.16.4"

idmap uid = 5000-50000
idmap gid = 5000-50000

ldap ssl = [Off|start tls] #whichever you use
```

Advanced: Idapsam - smb.conf II

Set password for ldap admin dn

```
smbpasswd -w [password]
```

Restart samba daemons

```
/etc/init.d/samba restart
```

Check log files and verify no errors so far

<http://www.samba.org/samba/docs/man/Samba-HOWTO-Collection/passdb.html>,
http://wiki.ubuntuusers.de/OpenLDAP_ab_Precise,
<https://help.ubuntu.com/12.04/serverguide/samba-ldap.html>, man smb.conf

Advanced: Idapsam - smbldap-tools I

Install smbldap-tools

```
apt-get install smbldap-tools #Bugged!  
apt-get source smbldap-tools  
./configure && make
```

Configure smbldap-tools

```
./smbldap-config.cmd  
#Leave values at default if not known
```

/etc/smbldap-tools/smbldap_bind.conf

```
slavePw="myadminpw" #Check!  
masterPw="myadminpw" #Check!
```

Advanced: Idapsam - smbldap-tools II

Use `smbldap-tools` to administrate accounts

```
smbldap-populate #Initialize everything
smbldap-useradd -a -P jonas2
    #New Unix and Samba user
smbldap-passwd jonas2
smbldap-userinfo -l jonas2
```

<http://www.samba.org/samba/docs/man/Samba-HOWTO-Collection/passdb.html>,
http://wiki.ubuntuusers.de/OpenLDAP_ab_Precise,
<https://help.ubuntu.com/12.04/serverguide/samba-ldap.html>, `man smbldap-*`, `man smb.conf`

Layout

NFS

Samba

SSHFS

Sun Grid Engine

SSHFS

As user root

```
apt-get install sshfs  
adduser jonas fuse
```

As user jonas (after re-login)

```
mkdir ~/L0home  
sshfs jonas@i12r-tbl.informatik.  
tu-muenchen.de:/home/tbl2013/jonas  
~/L0home -p 8574  
fusermount -u ~/L0home #for unmounting
```

- ▶ FUSE is the user-space filesystem, giving non-root users the ability to run their own filesystems

Layout

NFS

Samba

SSHFS

Sun Grid Engine

Installation

```
apt-get install gridengine-{master,exec,  
qmon,client}
```

Configure SGE automatically?

<Yes>

SGE cell name:

jonas

<Ok>

SGE master hostname:

jonas.tbl

<Ok>

qmon - Overview

User Configuration

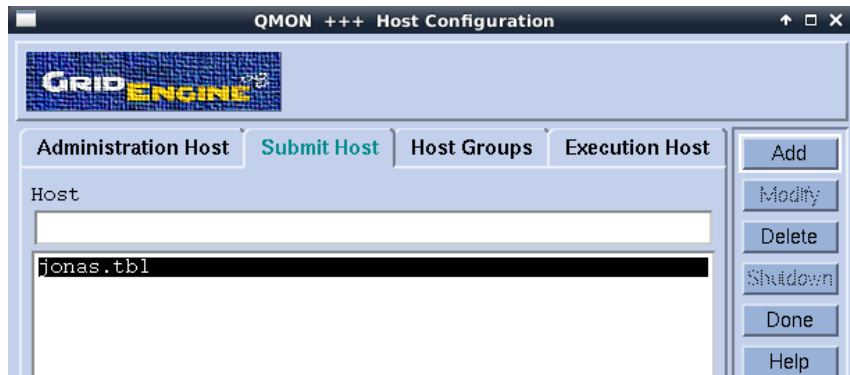
Host Configuration

Queue Control



Policy Configuration

qmon - Host configuration



qmon - Queue Control I

Queue Configuration - Add

GRIDENGINE

Queue Name: default

Hostlist: jonas.tbl


New Host/Hostgroup:

Attributes for Host/Hostgroup: a/

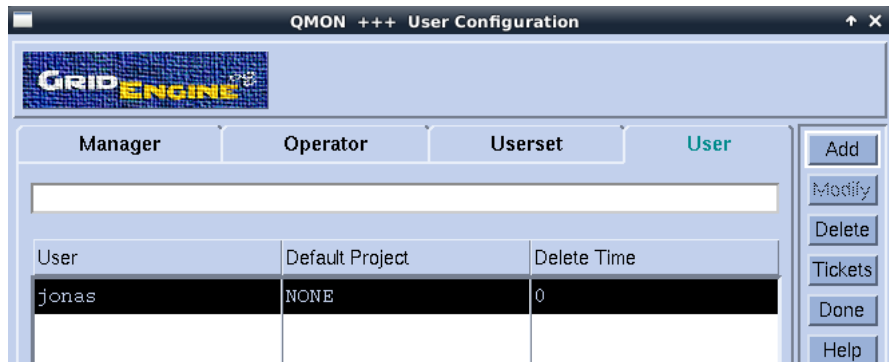
New Host/Hostgroup:

User Access	Project Access	Owners
Load/Suspend Thresholds	Limits	Complexes
General Configuration	Execution Method	Checkpointing
Sequence Nr	0	Calendar
Processors	UNDEFINED	Notify Time
tmp Directory	/tmp	Job's Nice
Shell	/bin/sh	Slots
Shell Start Mode	posix_compliant	Type
Initial State	default	Batch
Rerun Jobs	<input type="checkbox"/>	Interactive

qmon - Queue Control II

QMON +++ Cluster Queues								
								
Cluster Queues		Queue Instances				Hosts		
CLUSTER QUEUE	CQLOA	USED	RES	AVAIL	TOTAL	aoACDS	cdsuE	s
default	0.21	0	0	1	1	0	0	0

qmon - User configuration



QMON +++ User Configuration

GRID ENGINE v3

Manager Operator Userset **User**

Search bar

User	Default Project	Delete Time
jonas	NONE	0

Add
Modify
Delete
Tickets
Done
Help

qmon - Policy configuration

@ { QMON }

GRIDENGINE

Priority: 0 Urgency: 0.5 Ticket: 0.5

Weight Deadline: 3.6e+06

Weight Waiting Time: 0.278

Current Active Tickets: 0 Edit Tickets

Share Tree Tickets: 0 Total Share Tree Tickets: 10000

Functional Tickets: 0 Total Functional Tickets: 0

Override Tickets: 0

☐ Share Override Tickets

☐ Share Functional Shares

☐ Report Pending Job Tickets

Maximum Functional Jobs: 200

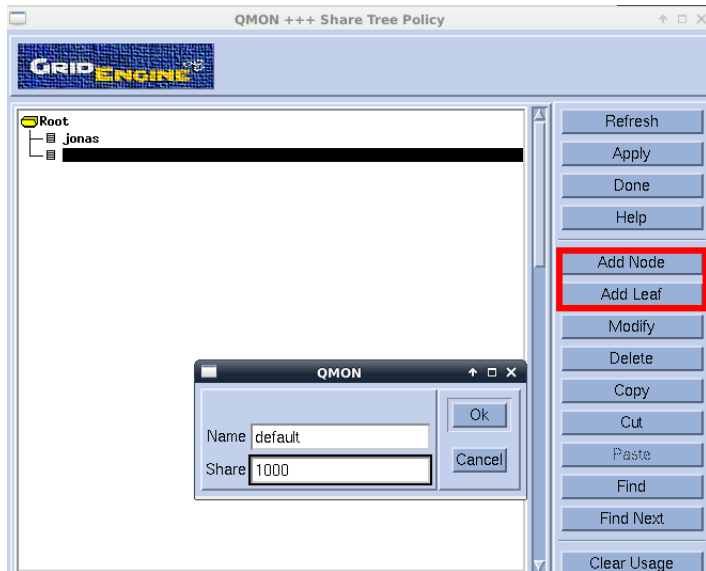
Maximum Pending Tasks Per Job: 50

Ticket Policy Hierarchy: OFS

Users Allowed To Submit Deadline Jobs **Share Tree Policy** Functional Policy Override Policy

Refresh Apply Done Help

qmon - Share Tree Policy



Job submission I

```
mkdir /home/jonas/sgeout  
chmod -R a+rwx /home/jonas/sgeout
```

myjob.sh

```
#!/bin/bash  
#$ -j y  
#$ -o /home/jonas/sgeout/  
  
sleep 5  
/bin/date
```

Job submission II

```
jonas@jonas:~$ qsub myjob.sh
Your job 7 ("myjob.sh") has been submitted
jonas@jonas:~$ qstat
```

job-ID	prior	name	user	state	submit/start at	queue
7	0.75000	myjob.sh	jonas	r	06/22/2013 14:23:24	default@jonas.tbl

```
jonas@jonas:~$ ls -la /home/jonas/sgeout
total 24
drwxrwxrwx 2 root root 4096 Jun 22 14:23 .
drwxr-xr-x 41 jonas 1000 4096 Jun 22 14:21 ..
-rw-r--r-- 1 jonas jonas 30 Jun 22 13:17 myjob.sh.o3
-rw-r--r-- 1 jonas jonas 30 Jun 22 14:22 myjob.sh.o5
-rw-r--r-- 1 jonas jonas 30 Jun 22 14:22 myjob.sh.o6
-rw-r--r-- 1 jonas jonas 30 Jun 22 14:23 myjob.sh.o7
jonas@jonas:~$ cat /home/jonas/sgeout/myjob.sh.o7
Sat Jun 22 14:23:29 CEST 2013
```

Administration / Diagnosis

```
qstat -j [jobid]  
qstat -explain [error states]  
qdel [jobid]
```

Advanced: L2 as Execution Host; On L2 I

```
apt-get install gridengine-{client,exec}
```

Choose automatic, cell name `jonas` and master
hostname `jonas.tbl`

L1 doesn't know about new execution host yet, we'll
fix that later

```
root@debian:/home/jonas12# cat /tmp/execd_messages.3334
06/22/2013 10:23:12|  main|debian|E|got send error
06/22/2013 10:23:12|  main|debian|E|can't get configuration from qmaster -- back
grounding
```

Advanced: L2 as Execution Host; On L2 II

/etc/hosts

```
127.0.0.1      jonasl2.tbl    jonasl2
192.168.16.34 jonasl2.tbl
```

/etc/hostname

```
jonasl2.tbl
```

Restart L2

Mount L1 home so jobs can access it

```
mount -t nfs4 192.168.16.4:/home /home/
```

<http://verahill.blogspot.de/2012/06/setting-up-sun-grid-engine-with-three.html>

Advanced: L2 as Execution Host; On L1 I

/etc/hosts

```
192.168.16.34  jonasl2.tbl
```

Add exec host: qmon -> Host configuration

The screenshot shows the qmon interface with the 'Execution Host' tab selected. A red arrow points from the 'Add' button in the top right to the 'Add/Modify Exec Host' dialog box. The dialog box has a 'Host' field containing 'jonasl2.tbl'. The background window shows a list of hosts with 'global' and 'jonas.tbl' listed under the 'Hosts' column.

Administration Host	Submit Host	Host Groups	Execution Host
Hosts			
Load Scaling			
global			
jonas.tbl			

Add/Modify Exec Host

Host:

Fixed Attributes | User Access | Reporting Variables | Project Access

Buttons: Add, Modify, Delete, Ok, Cancel

Advanced: L2 as Execution Host; On L1 II

Modify the queue: qmon -> Queue control

The screenshot shows the QMON Cluster Queues interface. The main window displays a table of Cluster Queues. The 'default' queue is selected, and the 'Modify default' dialog is open. A red arrow points to the 'Modify' button in the right-hand menu.

QMON +++ Cluster Queues

Cluster Queues

CLUSTER QUEUE	CQLOA	USED	RES	AVAIL	TOTAL	aoACDS	cdsuE	s
default	0.27	0	0	1	1	0	0	0

Queue Instances

Hosts

Modify default

GRID ENGINE

Modify default

Queue Name: default

Hostlist

- jonas.tbl
- jonas12.tbl

New Host/Hostgroup

Advanced: L2 as Execution Host; On L1

III

Check that execution host is correctly recognized

```
root@jonas:/home/jonas# qghost
HOSTNAME                               ARCH          NCPU  LOAD  MEMTOT  MEMUSE  SWAPTO  SWAPUS
-----
global                                -             -     -     -        -        -        -
jonas.tbl                             1x26-amd64    1  0.22   1.1G    757.9M   383.0M   2.9M
jonas12.tbl                           1x26-amd64    1  0.01   497.0M   25.8M    94.0M    0.0
root@jonas:/home/jonas# qstat -f
queueName                               qtype resv/used/tot. load_avg arch          states
-----
default@jonas.tbl                       BIP    0/0/1           0.22    1x26-amd64
default@jonas12.tbl                     BIP    0/0/1           0.01    1x26-amd64
```

<http://verahill.blogspot.de/2012/06/setting-up-sun-grid-engine-with-three.html>

Advanced: L2 as Execution Host; On L2: Extras I

Ensuring users are present at all execution hosts

```
apt-get install nscd libpam-ldap  
libnss-ldap ldap-utils
```

Use same configuration as for `dc=tbl` on L1 but give 192.168.16.4 as `ldap-server` address

Advanced: L2 as Execution Host; On L2: Extras II

`/etc/libnss-ldap.conf` and `/etc/pam_ldap.conf`

```
host 192.168.16.4
```

```
#uri ldapi:///192.168.16.4
```

```
pam-auth-upgrade
```

Try to login with an ldap defined user

Advanced: L2 as Execution Host; Success

```
jonas@jonas:~$ qsub myjob.sh
Your job 41 ("myjob.sh") has been submitted
jonas@jonas:~$ qsub myjob.sh
Your job 42 ("myjob.sh") has been submitted
jonas@jonas:~$ qsub myjob.sh
Your job 43 ("myjob.sh") has been submitted
jonas@jonas:~$ qstat -f
```

queue	name	qtype	resv/used/tot.	load_avg	arch	states
default	jonas.tbl	BIP	0/1/1	0.05	1x26-amd64	
42	0.75000 myjob.sh	jonas	r	06/22/2013 21:22:41	1	
default	jonasl2.tbl	BIP	0/1/1	0.01	1x26-amd64	
41	1.00000 myjob.sh	jonas	r	06/22/2013 21:22:41	1	
#####						
- PENDING JOBS - PENDING JOBS - PENDING JOBS - PENDING JOBS - PENDING JOBS						
#####						
43	0.16667 myjob.sh	jonas	qw	06/22/2013 21:22:32	1	

Thank you

References I

NFS

- ▶ `man exports, mount, fstab`
- ▶ http://www.centos.org/docs/5/html/Deployment_Guide-en-US/s1-nfs-server-config-exports.html
- ▶ <http://nfs.sourceforge.net/>

Samba

- ▶ `man smbpasswd, smb.conf, mount.cifs, smbldap-*`
- ▶ <https://wiki.archlinux.org/index.php/Thunar>
- ▶ <http://www.samba.org/samba/docs/man/Samba-HOWTO-Collection/passdb.html>

References II

- ▶ http://wiki.ubuntuusers.de/OpenLDAP_ab_Precise
- ▶ <https://help.ubuntu.com/12.04/serverguide/samba-ldap.html>

SSHFS

- ▶ `man sshfs, fusermount`

Sun Grid Engine

- ▶ `man qsub, qstat, qdel`
- ▶ <http://verahill.blogspot.de/2012/06/setting-up-sun-grid-engine-with-three.html>